

The Architecture of Business Intelligence

Adapted from the forthcoming book
Competing on Analytics: The New Science of Winning
Harvard Business School Press, March 2007

by **Thomas H. Davenport and Jeanne G. Harris**
prepared for Accenture Information Management Services



**>
accenture**

Institute for
High Performance
Business

March 2007

• Consulting • Technology • Outsourcing

Many companies today are collecting and storing a mind-boggling quantity of data. The numbers are hard to fathom: in just a few years, the common terminology for data volumes has grown from megabytes to gigabytes to terabytes (a trillion bytes). The size of some corporate databases is even approaching one petabyte (a quadrillion bytes). Just how much data is this? To put the 583 TB in Wal-Mart's databases into perspective, for example, consider that in 2006, U.S. Library of Congress's print collection was roughly 20 TB. And gargantuan storage is not the only technological frontier: high-end 64-bit processors and specialty "data appliances" can quickly churn through virtually unfathomable amounts of data.

However, while organizations have more data than ever at their disposal, they rarely impose sufficient order on it and thus get limited value from all that information. Further, many IT departments lack the capabilities to do more than support and maintain basic transactional and reporting capabilities. In short, while improvements in technology's ability to store data have been astonishing, most organizations struggle to manage, analyze and apply it.

Companies such as Netflix, Tesco, the global cement manufacturer CEMEX, Harrah's Entertainment and Capital One have only two things in common: they compete on the basis of their analytical capabilities and they are highly successful. By using business intelligence to make better decisions and to extract maximum value from their business processes, these organizations are able to identify their most profitable customers, accelerate product innovation, optimize supply chains and pricing, and identify the true drivers of financial performance.

Building a robust analytical capability requires much more than just collecting and storing data in large quantities. There are many moving pieces to put in place, including software applications, technology, data, processes, metrics, incentives, skills, culture, and sponsorship. An important initial step for any organization is to understand the elements of a business intelligence architecture so that structure can be imposed on otherwise chaotic bytes of data.

What is business intelligence?

To make sure the IT environment fully addresses an organization's needs at each stage of analytical competition, companies must incorporate analytics and other business intelligence technologies into their overall IT architecture.

Technologists use the term "business intelligence" (often shortened as BI) to encompass analytics as well as the processes and technologies used for collecting, managing, and reporting decision-oriented data. The business intelligence architecture (a subset of the overall IT architecture) is an umbrella term for an enterprise-wide set of systems, applications, and governance processes that enable sophisticated analytics, by allowing data, content and analyses to flow to those who need them, when they need them.

Top management, functional heads, knowledge workers, and statisticians all need information of these types at various times and in various forms. The BI architecture must be able to quickly provide users with reliable, accurate information and help them make decisions of widely varying complexity (See "Top Ten Signs of Effective Business Intelligence"). It also must make information available through a variety of distribution channels, including traditional reports, ad hoc analysis tools, corporate dashboards, spreadsheets, e-mail, and pager alerts. This task is often daunting: Amazon.com,

for example, spent more than ten years and a billion dollars building, organizing and protecting its data warehouses.

Complying with legal and regulatory reporting requirements is another activity that depends on a robust BI architecture. The Sarbanes-Oxley Act of 2002, for example, requires executives, auditors and other users of corporate data to demonstrate that their decisions are based on trustworthy, meaningful, authoritative and accurate data. It also requires them to attest that the data provides a clear picture of the business, major trends, risks and opportunities.

Responsibility for getting the data, technology and processes right is the job of the IT architect. This executive (working closely with the CIO) must determine how the components of the IT infrastructure as a whole (hardware, software and networks) will work together to provide the data, technology and support needed by the business. In large established organizations the IT infrastructure can sometimes appear to have been constructed in a series of weekend handyman jobs. It does the job it was designed to do but is apt to create problems whenever it is applied to another purpose.

Breaking the business intelligence architecture into its six elements can help IT executives leverage the analytical power of their IT investment.

1. Data management that defines how the right data is acquired and managed.
2. Transformation tools and processes that describe how the data is extracted, cleaned, transmitted and loaded to "populate" databases.
3. Repositories that organize data and metadata (information about the data) and store it for use.

4. Applications and other software tools used for analysis.

5. Presentation tools and applications that address how information workers and non-IT analysts will access, display, visualize and manipulate data.

6. Operational processes that address how important administrative activities such as security, error handling, "auditability," archiving and privacy are resolved.

We'll look at each element in turn, with particular attention to data management since it drives all the other architectural decisions.

Business intelligence, analytics and high performance

As part of Accenture's research into high-performance businesses, we have found that a growing number of companies have recognized the power of leveraging data-driven insights through the use of business intelligence. Some forward-thinking companies have gone a step further and are building their competitive strategies around analytics—that is, the extensive use of data, statistical and quantitative analysis, explanatory and predictive models, and fact-based management to drive decisions and actions. Accenture's high performance business research found a powerful link between organizations with pronounced analytical orientations and market out-performance.

Top Ten Signs of Effective Business Intelligence

1. Managers and analysts have direct, nearly instantaneous access to data; they never argue over whose numbers are accurate.
2. Information workers spend their time analyzing data and understanding its implications rather than collecting and formatting data.
3. Managers focus on improving processes and business performance, not culling data from laptops, reports, and transaction systems.
4. A hypothesis can be quickly analyzed and tested without a lot of manual behind-the-scenes preparation.
5. Data is managed from an enterprise-wide perspective throughout its life cycle, from its initial creation to archival or destruction.
6. Rather than have data warehouse or business intelligence initiatives, companies manage data as a strategic corporate resource in all business initiatives.
7. Both the supply and demand sides of the business rely on forecasts that are aligned and have been developed using a consistent set of data.
8. High-volume, mission-critical decision-making processes are highly automated and integrated.
9. Data is routinely and automatically shared between the company and its customers and suppliers.
10. Reports and analyses seamlessly integrate and synthesize information from many sources.

Data management

The goal of a well-designed data management strategy is to ensure that an organization has the right information and uses it appropriately. Large companies invest millions of dollars in systems that snatch data from every conceivable source. Systems for enterprise resource planning, customer relationship management, and point-of-sale transactions, among others, ensure that no transaction or exchange occurs without leaving a mark. Many organizations also purchase externally gathered data from syndicated providers such as IRI and AC Nielsen in consumer products and IMS Health in pharmaceuticals.

In this environment, data overload can be a real problem for time-stressed executives. But the greatest data challenge facing companies is “dirty” data: information that is inconsistent, fragmented and out of context. Even the best companies often struggle to address their data issues. We found that companies that compete on analytics devote extraordinary attention to data management processes and governance. Capital One, for example, estimates that 25 percent of its IT organization works on data issues—an unusually high percentage compared with other firms.

There's a significant payoff, however, awaiting companies that invest the effort to master data management. For example, Continental Airlines integrates ten TB (terabytes) of data from 25 operational systems into its data warehouse. The data is used in analytical applications for both real-time alerts

and long-range strategic analysis. Alerts notify customer agents of delays in incoming flights and identify incoming frequent-flier customers who are assigned alternative flights if they are unlikely to make their connections. Marketing analysts use other data collected by the systems to study customer and pricing trends, and logistical analysts plan the optimal positioning of planes and crews. The company estimates that it has saved more than \$250 million in the first five years of its data warehousing and business intelligence activities—representing an ROI of more than 1,000 percent.

To achieve the benefits of analytical competition, IT and business experts must tackle their data issues by answering the following five questions:

1. What data is needed to compete on analytics? The question behind this question is, What data is most valuable for competitive differentiation and business performance?

Ensuring that analysts have access to the right data can be difficult. Sometimes a number is needed: the advent of credit scores made the mortgage-lending business more efficient by replacing qualitative assessments of consumer creditworthiness with a single, comparative metric. But not everything is reducible to a number. An employee's performance rating does not give as complete a picture of his work over a year as a manager's written assessment.

The situation is complicated when business and IT people blame each other when the wrong data is collected or the right data is not available. Studies show that IT executives believe business managers do not understand what data they need and, conversely, that business managers believe IT executives lack the business acumen to make meaningful data available. While there is no easy solution to this problem, the beginning of one is for business leaders and IT managers to pledge to work together.

The importance of collaboration between business and IT leaders can be seen in the complexity an insurer faces in defining health care customers. An insurance company has many different customers: corporations that contract for policies on behalf of their employees, individual subscribers, and members of the subscribers' families. Each individual has a unique medical history and current needs. Some may be covered by other providers, including the government. The insurance company also has relationships with service providers such as hospitals, HMOs and doctors, which come in their own varieties and are connected in many combinations.

Companies must think through such difficult questions before they can make full use of business intelligence capabilities.

2. Where can this data be obtained?

Data for business intelligence originates from many places, but it must be managed through an enterprise-wide infrastructure. Only by this means will it be streamlined, consistent and scalable. Having common applications and data across the enterprise is critical because it helps yield a "single version of the truth."

To organize internal information, an organization's enterprise systems are a logical starting point. Enterprise systems—integrated software applications that automate, connect and manage information flows for business processes such as order fulfillment—often help companies move along the path toward analytical competition: they provide consistent, accurate and timely data for such tasks as financial reporting and supply chain optimization. Vendors

increasingly are embedding analytical capabilities into their enterprise systems so that users can develop sales forecasts and model alternative solutions to business problems.

In addition to corporate systems, an organization's personal computers and servers are loaded with data. Databases, spreadsheets, presentations and reports are all sources of data. Sometimes these sources are stored in a common knowledge management application, but they are often not available across the entire organization.

For external information, managers can purchase data from firms that provide financial and market information, consumer credit data and market measurement. Governments at all levels are some of the biggest information providers, and company Web sites are another powerful resource. Data can also come from the sources popularized in the movies: e-mail, voice applications, images (maps and photos available through the Internet) and biometrics (fingerprints and iris identification). The further the data type is from standard numbers and letters, however, the harder it is to integrate with other data and analyze.

More data about the physical world, through sensor technology and radio frequency identification (RFID) tags, is becoming available as well. For example, a case of wine can be monitored to see whether it is being kept at the proper temperature.

It can be difficult and expensive to capture some highly valuable data. (In some cases, it might even be illegal—for example, sensitive customer information or competitor intelligence about new product plans or pricing strategies.) Analytical competitors adopt innovative approaches to gain permission to collect the data they need. Progressive Insurance's TripSense program offers discounts to customers who agree to install a device that collects data about their driving behavior. Former CEO Peter Lewis sees this capability as the key to more accurate pricing and capturing the most valuable customers: "It's about being able to charge them for whatever happens instead of what they [customers] say is happening. So what will happen? We'll get all the people who hardly ever drive, and our competitors will get stuck with the higher risks."

3. How much data is needed? In addition to gathering the right data, companies need to collect a lot of it in order to distill trends and predict behavior.

Two pitfalls must be balanced against this need for large quantities of data. First, companies have to resist the temptation to collect all the data they can "just in case." For one thing,

if executives have to wade through digital oceans of irrelevant data, they will give up before they drown. For another, data hoarders quickly learn that they cannot collect everything, and that even if they try, the costs outweigh the benefits.

Second, companies should avoid collecting data that is easy to capture but not necessarily important. Many IT executives advocate a low-hanging-fruit approach because it relieves them of responsibility for determining what information is valuable to the business. IT and business executives must be clear about what drives value in an organization; this understanding will prevent companies from collecting data indiscriminately.

4. How can we make data more valuable? Quantity without quality is a recipe for failure. Executives are aware of the problem: in a survey of the challenges organizations face in developing a business intelligence capability, data quality was second only to budget constraints. Organizations tend to store their data in hard-walled, functional silos. As a result, the data is generally a disorganized mess. For most organizations, differing definitions of key data elements such as customer or product add to the confusion. When Canadian Tire, for example, set out to create a structure for its data, it found that the company's data warehouse could yield as many as six different numbers for inventory levels. Other data was not available at all, such as comparison sales figures for certain products sold in its 450-plus stores throughout Canada. Over several years, the company created a plan to collect new data that fit the company's analytical needs.

Several characteristics increase the value of data:

It is correct. While some analyses can get by with ballpark figures and others need precision to several decimal points, all must be informed by data that passes the credibility tests of the people reviewing it.

It is complete. The definition of "complete" will vary according to whether a company is selling cement, credit cards or season tickets, but completeness will always be closely tied to an organization's distinctive capability—the integrated business processes and capabilities that together serve customers in ways that are differentiated from competitors and that create an organization's formula for business success.

It is current. Again, the definition of "current" may vary; for some business problems, such as a major medical emergency, data must be available instantly to deploy ambulances and emergency personnel (also known as "zero latency"); for most other business decisions, such as a budget forecast, it just needs to be updated periodically—daily, weekly or monthly.

It is consistent. In order to help decision makers end arguments over whose data is correct, standardization and common definitions must be applied to it. Eliminating redundant data reduces the chances of using inconsistent or expired data.

It is in context. When data is enriched with metadata (usually defined as “structured data about data”), its meaning and how it should be used become clear.

It is controlled. In order to comply with business, legal and regulatory requirements for safety, security, privacy and “auditability,” it must be strictly overseen.

5. What rules and processes are needed to manage the data from its acquisition through its retirement?

Each stage of the data management life cycle presents distinctive technical and management challenges.

Data acquisition. Creating or acquiring data is the first step. For internal information, IT managers should work closely with business process leaders. The goals include determining what data is needed and how to best integrate IT systems with business processes to capture good data.

Data cleansing. Between 25 and 30 percent of a BI initiative typically goes toward initial data cleansing: detecting and removing data that is out of date, incorrect, incomplete or redundant. IT’s role is to establish methods and systems to collect, organize, process and maintain information, but data cleansing is the responsibility of everyone who generates or uses data.

Data organization and storage. Once data has been acquired and cleansed, processes to systematically extract, integrate and synthesize it must be

established. The data must then be put into the right repository and format so that it is ready to use.

Data maintenance. After a repository is created and populated with data, managers must decide how and when the data will be updated. They must create procedures to ensure data privacy, security and integrity (protection from corruption or loss via human error, software virus or hardware crash). And policies and processes must also be developed to determine when and how data that is no longer needed will be saved, archived or retired. Some companies have estimated that they spend \$500,000 on maintenance for every \$1 million spent on developing new business intelligence technical capabilities.

Once an organization has addressed data management issues, it must determine the technologies and processes needed to capture, transform and load data into a data warehouse.

Transformation tools and processes

For data to become usable by managers, it must first go through a process known in IT-speak as ETL, for extract, transform and load. While extracting data from its source and loading it into a repository are fairly straightforward tasks, cleansing and transforming data are not.

The first step is to clean and validate data using business rules and data cleansing tools such as Trillium. (For example, a simple rule might be to have a full nine-digit ZIP code for all U.S. addresses.) Transformation procedures

then define the business logic that maps data from its source to its destination. Both business and IT managers must expend significant effort in order to transform data into usable information. While automated tools from vendors such as Informatica, Ab Initio and Ascential can ease this process, considerable manual effort is still required. Informatica’s CEO, Sohaib Abbasi, estimates that “for every dollar spent on integration technology, around seven to eight dollars is spent on labor [for manual data coding].”

Transformation also entails standardizing data definitions to make certain that business concepts have consistent, comparable definitions across the organization. For example, a “customer” may be defined as a company in one system but as an individual placing an order in another. It also requires managers to decide what to do about data that is missing. These mundane but critical tasks require an ongoing effort, because new issues seem to constantly arise.

Data repositories

Organizations have several options for organizing and storing their analytical data:

Data warehouses are regularly updated databases that contain integrated data from different sources. They contain time-series (historical) data to facilitate the analysis of business performance over time. A data warehouse may be a module of an enterprise system or an

independent database. Some companies also employ a “staging” database that is used to get data from many different sources ready for the data warehouse.

A **data mart** can refer to a separate repository or to a partitioned section of a complete data warehouse. Data marts are generally used to support a single business function or process and usually contain some predetermined analyses so that managers without statistical expertise can slice and dice some data. Data marts can result in the balkanization of data and should only be used if the designers are confident that no broader set of data will ever be needed for analysis.

A **metadata repository** contains technical information and a data definition, including information about the source, how it is calculated, bibliographic information and the unit of measurement. It may include information about data reliability, accuracy and instructions on how the data should be applied. A common metadata repository used by all analytical applications is critical to ensure data consistency. Consolidating all the information needed for data cleansing into a single repository significantly reduces the time needed for maintenance.

Once the data is organized and ready, it is time to determine the analytical technologies and applications needed.

Analytical tools and applications

Choosing the right software tools or applications depends on several factors. The first task is to determine how thoroughly decision making should be embedded into business processes. Should a decision be automated or made by a person? If it is automated, technologies exist that both structure the workflow and provide decision rules—either quantitative or qualitative—to make the decision.

Next, companies must decide whether to use a third-party application or create a custom solution. The “make or buy” decision hinges upon whether a packaged solution exists and whether the level of skill required exists within the organization. A growing number of functionally or industry-specific business applications, such as capital-budgeting or mortgage-pricing models, now exist. Vendors of enterprise systems such as Oracle and SAP are building more analytical applications into their products. According to IDC, projects that implement a packaged analytical application yield a median ROI of 140 percent, while custom development using analytical tools yields a median ROI of 104 percent.

Nevertheless, powerful tools have been created that allow organizations to develop their own analyses (see “Analytical technologies”). Companies such as Business Objects and SAS offer product suites consisting of integrated tools and applications. Some tools are designed to slice and dice or to drill down to predetermined views of the data, while others are more statistically sophisticated. Some tools can accommodate a variety of data types, while

others are more limited (to highly structured data or textual analysis, for example). Some tools extrapolate from historical data, while others are intended to seek out new trends or relationships.

Whether a custom solution or off-the-shelf application is used, the business IT organization must accommodate a variety of tools for different types of data analysis.

Employees naturally tend to prefer familiar products, such as spreadsheets, even if they are ill suited for the analysis to be done. Another problem is proliferation of technologies. In a 2005 survey, respondents from large organizations reported that their organizations had, on average, 13 business intelligence tools from an average of 3.2 vendors.

In the past, different vendors had different capabilities—one might focus on financial reporting, another on ad hoc query, and yet another on statistical analysis. Today, however, leading providers have begun to offer business intelligence suites with stronger, more integrated capabilities.

Analytical technologies

Executives in organizations that are planning to become analytical competitors should be familiar with the key categories of analytical software tools:

Spreadsheets such as Microsoft Excel are the most commonly used analytical tools because they are easy to use and reflect users' mental models. Managers and analysts use them for "the last mile" of analytics—the stage right before the data is presented in report or graphical form for decision makers. But too many users attempt to use spreadsheets for tasks for which they are ill suited, leading to errors or incorrect conclusions. Spreadsheet programs generally provide the ability to manage data arrays across three dimensions (down, across, and worksheet pages), while OLAP models (see below) can have seven or more, and are therefore suited to more complex problems. Even when used properly, spreadsheets are prone to human error; more than 20 percent of spreadsheets have errors, and as many as 5 percent of all calculated cells are incorrect.

Online analytical processors are generally known by their abbreviation, OLAP, and are used for semi-structured decisions and analyses. While a relational database (or RDBMS)—in which data is stored in related tables—is a highly efficient way to organize data for transaction systems, it is not particularly efficient when it comes to analyzing array-based data (data that is arranged in cells like a spreadsheet), such as time series. OLAP tools are specifically designed for multidimensional, array-based problems. They organize data in

"data cubes" to enable analysis across time, geography, product lines and so on. Data cubes are simply collections of data in three variables or more that are prepackaged for reporting and analysis; they can be thought of as multidimensional spreadsheets. Unlike traditional spreadsheets, OLAP tools must deal with data proliferation, or the models quickly become unwieldy. For complex queries, OLAP tools are reputed to produce an answer in around 0.1 percent of the time that it would take for the same query to be answered using relational data. Business Objects and Cognos are among the leading vendors in this category.

Statistical or quantitative algorithms process quantitative data to arrive at an optimal target such as a price or a loan amount. In the 1970s, companies such as SAS and SPSS introduced packaged computer applications that made statistics much more accessible. Statistical algorithms also encompass predictive modeling applications, optimization and simulations.

Rule engines process a series of business rules that use conditional statements to address logical questions—for example, "If the applicant for a motorcycle insurance policy is male and under 25, and does not either own his own home or have a graduate degree, do not issue a policy." Rule engines can be part of a larger automated application or provide recommendations to users who need to make a particular type of decision. Fair Isaac, ILOG and Pegasystems are some of the major providers of rule engines for businesses.

Data mining tools draw on techniques ranging from arithmetic computation to artificial intelligence, statistics, decision trees, neural networks and Bayesian network theory. Their objective is to identify patterns in complex and ill-defined data sets. Sprint, for example, uses neural analytical technology to predict which customers are likely to switch wireless carriers and take their existing phone numbers with them. SAS offers both data and text mining capabilities and is a major vendor in both categories.

Text mining tools can help managers quickly identify emerging trends in near-real time. A "spider" (or data crawler) that identifies and counts words and phrases on Web sites, is an example. Text mining tools can be invaluable in sniffing out new trends or relationships. For example, by monitoring technical-user blogs, a vendor can recognize that a new product has a defect within hours of being shipped instead of having to wait for complaints to arrive from customers.

Simulation tools model business processes with a set of symbolic, mathematical, scientific, engineering and financial functions. Much as computer-aided design systems are used by engineers to model the design of a new product, simulation tools are used in engineering, R&D and a surprising number of other applications. For

example, simulations can be used to help users understand the implications of a change to a business process. They can also be used to help streamline the flow of information or products—for example, to help employees of health care organizations decide where to send donated organs according to such criteria as blood type and geography.

Emerging analytical technologies

These are some of the leading-edge technologies that will play a role in business intelligence over the next few years:

Text categorization is the process of using statistical models or rules to rate a document's relevance to a certain

topic. For example, text categorization can be used to dynamically evaluate competitors' product assortments on their Web sites.

Genetic algorithms are a class of stochastic optimization methods that use principles found in natural genetic reproduction (crossover or mutations of DNA structures). One common application is to optimize delivery routes.

Expert systems are not a new technology but one that is finally coming of age. Specialized artificial intelligence applications are capable of making expertise available to decision makers—imagine a “Warren Buffett in a box” that gives advice on investment decisions in changing market conditions.

Audio and video mining are much like text or data mining tools, but they look for patterns in audio or images, particularly full-motion images and sound.

Swarm intelligence, as observed in the complex societies of ants and bees, is technology used to increase the realism of simulations and to understand how low-level changes in a system can have dramatic effects.

Information extraction culls and tags concepts such as names, geographical entities and relationships from largely unstructured (and usually) textual data.

Presentation tools and applications

Business intelligence will only work if people can impart their insights to others through reporting tools, scorecards and portals. Presentation tools should allow users to create ad hoc reports, to interactively visualize complex data, to be alerted to exceptions through communication tools (such as e-mail, PDAs, or pagers), and to collaboratively share data. (Business intelligence vendors such as Business Objects, Cognos, SAS, and Hyperion sell product suites that include data presentation and reporting solutions. As enterprise systems become more analytical, vendors such as SAP and Oracle are rapidly incorporating these capabilities as well.)

Commercially purchased analytical applications usually have an interface to be used by information workers, managers

and analysts. But for proprietary analyses, the presentation tools determine how different classes of individuals can use the data. For example, a statistician could directly access a statistical model, but most managers would hesitate to do so.

A new generation of analytical tools—from vendors such as Spotfire, Visual Sciences and SAS—allow the manipulation of data and analyses through an intuitive visual interface. A manager, for example, could look at a plot of data, exclude outlier values, and compute a regression line that fits the data—all without any statistical skills. Because they permit exploration of the data without the risk of accidentally modifying the underlying model, visual analytics tools increase the number of users who can employ sophisticated analyses. At Vertex Pharmaceuticals, for example, CIO Steve Schmidt estimates that only 5 percent of his users can make effective use of algorithmic tools, but another 15 percent can manipulate visual analytics.

Operational processes

This element of the BI architecture answers questions about how the organization creates, manages and maintains data and applications. It details how a standard set of approved tools and technologies can be used to ensure the reliability, scalability and security of the IT environment. Standards, policies and processes must also be defined and enforced across the entire organization.

Issues such as privacy and security as well as the ability to archive and audit the data are of critical importance to data integrity. This is a business as well as a technical concern, because lapses in privacy and security (for example, if customer credit card data is stolen) can have dire consequences. Executives can be found criminally negligent if they fail to establish procedures to document and demonstrate the validity of data used for business decisions.

Conclusion

Top management can help the IT architecture team plan a robust technical environment by establishing guiding principles for analytical architecture. Those principles can help to ensure that architectural decisions are aligned with business strategy, corporate culture and management style.

To make that happen, senior management must be committed to the process. Working with IT, senior managers must establish and rigorously enforce comprehensive data management policies, including data standards and consistency in data definitions. They must be committed to the creation and use of high-quality data that is scalable, integrated, well documented, consistent and standards-based. And they must emphasize that the business intelligence architecture should be flexible and able to adapt to changing business needs and objectives. A rigid architecture will not serve the needs of the business in a fast-changing environment.

About the authors

Jeanne G. Harris (jeanne.g.harris@accenture.com) is a senior executive research fellow at the Institute for High Performance Business in Wellesley, Massachusetts.

Thomas H. Davenport is the former director of the Accenture Institute for High Performance Business. He holds the President's Chair in Information Technology and Management at Babson College in Wellesley, Massachusetts.

About Accenture

Accenture is a global management consulting, technology services and outsourcing company. Committed to delivering innovation, Accenture collaborates with its clients to help them become high performance businesses and governments. With deep industry and business process expertise, broad global resources and a proven track record, Accenture can mobilize the right people, skills and technologies to help clients improve their performance. With approximately 146,000 people in 49 countries, the company generated net revenues of US\$16.65 billion for the fiscal year ended August 31, 2006. Accenture's home page is www.accenture.com.

About the Accenture Institute for High Performance Business

The Accenture Institute for High Performance Business creates strategic insights into key management issues through original research and analysis. Its management researchers combine world-class reputations with Accenture's extensive consulting, technology and outsourcing experience to conduct innovative research and analysis into how organizations become and remain high-performance businesses.



High performance. Delivered.

Copyright © 2007 Accenture
All rights reserved.

Accenture, its logo, and
High Performance Delivered
are trademarks of Accenture.